arXiv:1611.06983v1 [math.CO] 21 Nov 2016

DIAMETER AND AUTOMORPHISMS OF GELFAND-TSETLIN POLYTOPES

YIBO GAO, BENJAMIN KRAKOFF, AND LISA YANG

Abstract. We determine the diameter of the 1-skeleton and the combinatorial automorphism group of any Gelfand-Tsetlin polytope GT_λ associated to an integer partition λ.

1. Introduction and Statement of Results

Gelfand-Tsetlin (GT) polytopes are compact convex polytopes defined by a set of linear inequalities depending on a partition λ as shown in Figure 1. The polytope GT_λ corresponds to all fillings of this triangular array with real numbers such that all rows and columns are weakly increasing.

$$
\begin{array}{ccccccccc}
\lambda_1 & & & & & & & & \\
\wedge & & & & & & & & \\
x_{2,1} & \leq & \lambda_2 & & & & & & \\
\wedge & & \wedge & & & & & & \\
x_{3,1} & \leq & x_{3,2} & \leq & \lambda_3 & & & & \\
\wedge & & \wedge & & \wedge & & & & \\
x_{4,1} & \leq & x_{4,2} & \leq & x_{4,3} & \leq & \lambda_4 & & \\
\vdots & & \vdots & & \vdots & & \ddots & & \\
x_{n,1} & \leq & x_{n,2} & \leq & \cdots & \leq & x_{n,n-1} & \leq & \lambda_n
\end{array}
$$

Figure 1. Inequality constraints of GT polytopes.

GT polytopes arise from the study of representations of $GL_n(\mathbb{C})$ and have connections to areas of representation theory and algebraic geometry (see for example [KST12]). For any integer partition $\lambda = (\lambda_1, \ldots, \lambda_n)$, let n be the length of λ and let GT_λ denote the GT polytope associated to λ. Then the integral points within GT_λ are in bijection with the number of semi-standard Young tableaux of shape λ with tableaux entries bounded by n. Furthermore, the integral points of GT_λ parametrize a *Gelfand-Tsetlin basis* of the GL_n-module with highest weight λ, so the number of integral points equals the dimension of this module. GT polytopes can also be viewed as the marked order polytope of a poset which is discussed in [Sta86] and [ABS11].

This paper describes the diameter of the 1-skeleton and the combinatorial automorphism group of GT polytopes. Since our results are purely combinatorial, it suffices to consider

Date: November 22, 2016.

partitions of the form $\lambda = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m})$ where $a_i > 0$ for all $1 \le i \le m$. This is explained in Section 2 and in Remark 2.5.

The 1-skeleton of a polytope is the graph obtained by looking only at its vertices and edges (the 0 and 1 dimensional faces). The shortest path between two vertices is the path with the minimum number of edges. The diameter of a graph is the maximum over all pairs of vertices of the length of the shortest path between the pair of vertices. Our first result is a formula for the diameter of the 1-skeleton of GT_λ which we denote by $\mathrm{diam}(\mathrm{GT}_\lambda)$.

We adopt the following notational conventions:

- The Kronecker delta function is defined as $\delta_{x,y} := \begin{cases} 1, & \text{if } x = y, \\ 0, & \text{if } x \neq y. \end{cases}$
- For any $a \in \mathbb{Z} > 0$, let S_a denote the symmetric group on a set of size a.
- For $\lambda = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m})$, its reverse partition is $\lambda' := (1^{a_m}, 2^{a_{m-1}}, \ldots, m^{a_1})$. If $a_i = a_{m+1-i}$ for all $1 \le i \le m$, then we say $\lambda = \lambda'$ and call λ a *reverse symmetric* partition.

Theorem 1.1 (Diameter of 1-skeleton)**.** *For any GT polytope GT_λ, $diam(GT_\lambda) = 2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}$.*

A polytope's face poset (or face lattice) is formed by ordering its faces by inclusion. A combinatorial automorphim of a polytope is an automorphism of its face poset. Our second result is a description of the combinatorial automorphism group of GT_λ.

Theorem 1.2 ($m = 2$ Automorphisms)**.** *Suppose $\lambda = (1^{a_1}, 2^{a_2})$ and $a_1, a_2 \ge 2$. If $a_1 = a_2 = 2$, then*

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2.$$

Otherwise,

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}},$$

where D_4 is the dihedral group of order 8 and $\mathbb{Z}_2$ is the cyclic group of order 2.

Theorem 1.3 ($m \ge 3$ Automorphisms)**.** *Suppose $\lambda = (1^{a_1}, \ldots, m^{a_m})$ and $m \ge 3$. Let r_1 be the number of k such that $a_k, a_{k+1} \ge 2$. Let $r_2 = 1$ if $\lambda = \lambda'$ and let $r_2 = 0$ otherwise. Then*

$$Aut(GT_\lambda) \cong (S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}) \ltimes_\varphi \mathbb{Z}_2^{r_2}$$

where if $r_2 = 1$, then $\varphi : \mathbb{Z}_2 \to Aut(S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1})$ sends the nonidentity element of $\mathbb{Z}_2$ to the map sending $(\sigma_1, \sigma_2, z_1, \ldots, z_{r_1}, z_{r_1+1}) \mapsto (\sigma_2, \sigma_1, z_{r_1}, \ldots, z_1, z_{r_1+1})$.

Example 1.4. Figure 2 shows the Gelfand-Tsetlin polytope $\mathrm{GT}_{(1,2,3)}$. It is clear that the diameter of $\mathrm{GT}_{(1,2,3)}$ is 2 and the automorphism group is $\mathbb{Z}_2 \times \mathbb{Z}_2$. One generator is given by rotating 180° about the axis containing vertex 4 and vertex 7. The other generator is given by fixing vertex 1 and interchanging vertex 2 with vertex 3, vertex 5 with vertex 6, and vertex 4 with vertex 7. These results are predicted by our theorems.

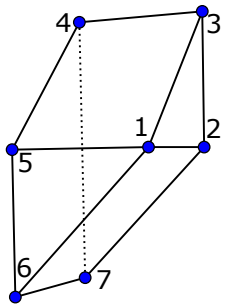

FIGURE 2. Polytope $\mathrm{GT}_{(1,2,3)}$.

In Section 2, we review background on GT polytopes including how to model their face poset using combinatorial objects called *ladder diagrams*. Then we present a proof of Theorem 1.1 in Section 3. In Section 4, we identify the generators of $\mathrm{Aut}(\mathrm{GT}_\lambda)$ and present proofs of Theorems 1.2 and 1.3.

2. Preliminaries

In this section, we formally define GT polytopes and describe how to view their faces as combinatorial objects called *ladder diagram*. These diagrams are the primary objects we will use to model the face lattice of GT_λ.

2.1. **Gelfand-Tsetlin Polytopes.** A partition of s is a sequence of weakly increasing positive integers $\lambda = (\lambda_1 \le \lambda_2 \le \ldots \le \lambda_n)$ such that $\sum_{i=1}^{n} \lambda_i = s$. We will often use multiplicative notation for λ and write $\lambda = (\lambda_1^{a_1}, \lambda_2^{a_2}, \ldots, \lambda_m^{a_m})$ for $a_1, \ldots, a_m \in \mathbb{Z}_{\ge 0}$ to denote a partition with a_1 copies of λ_1, a_2 copies of λ_2, and so forth. We may omit writing the term $\lambda_i^{a_i}$ if $a_i = 0$.

Definition 2.1 (GT Polytope)**.** Given a partition $\lambda = (\lambda_1, \ldots, \lambda_n)$, the Gelfand-Tsetlin Polytope GT_λ is the set of points $\vec{x} = (x_{i,j})_{1 \le j \le i \le n} \in \mathbb{R}^{n(n+1)/2}$ such that $x_{i,i} = \lambda_i$ for $1 \le i \le n$ and such that the following inequalities are satisfied:

(1) $x_{i-1,j} \le x_{i,j} \le x_{i+1,j}$,
(2) $x_{i,j-1} \le x_{i,j} \le x_{i,j+1}$.

Suppose for some $i < j$ that $\lambda_i = \lambda_j$. Then for every $i \le i', j' \le j$, we are forced to have $x_{i',j'} = \lambda_i$. Whenever such a situation occurs, we say that the coordinate $x_{i',j'}$ is *fixed*. In general, GT_λ will be a polytope in $\mathbb{R}^d$ where d is at most $n(n-1)/2$.

These constraints can be visualized in a triangular array as shown in Figure 1.

We adopt the following notational conventions:

- Let n denote the length of λ. We specify a partition as $\lambda = (\lambda_1, \ldots, \lambda_n)$. We typically only consider partitions of the form $\lambda = (1^{a_1}, \ldots, m^{a_m})$ where $a_i > 0$ for all $1 \le i \le m$.

- Let m denote the number of distinct values of λ. This agrees with our notation for partitions of the form $\lambda = (1^{a_1}, \ldots, m^{a_m})$ where $a_i > 0$ for all $1 \le i \le m$.
- Let d denote the dimension of GT_λ. Note that $d = \binom{n}{2} - \sum_{i=1}^{m} \binom{a_i}{2}$.
- Let $\mathcal{F}(\mathrm{GT}_\lambda)$ denote the face poset of GT_λ ordered by inclusion.
- Let $\mathcal{I}_n = \{(i,j) : 1 \le j \le i \le n\}$ denote the triangular grid with shape shown in Figure 1.

2.2. **Ladder Diagrams.** For every λ, we define a graph Γ_λ such that faces of GT_λ correspond to subgraphs of Γ_λ with certain restrictions. These subgraphs are the ladder diagrams introduced in [ACK16].

Let Q be the infinite graph corresponding to first quadrant of the Cartesian plane, i.e. let Q have vertices (i,j) for all $i,j \ge 0$ and edges $\{(i,j),(i+1,j)\}$ and $\{(i,j),(i,j+1)\}$. For the sake of convenience, define $a_0 := 0$ and $s_j := \sum_{i=0}^{j} a_i$ for $0 \le j \le m$.

Definition 2.2 (Γ_λ and Ladder Diagrams)**.** For $\lambda = (1^{a_1}, \ldots, m^{a_m})$, the grid Γ_λ is an induced subgraph of Q constructed as follows. We identify the vertex at $(0,0)$ as the *origin* and the vertices $t_j = (s_j, n - s_j)$ for $0 \le j \le m$ as *terminal vertices*. Γ_λ consists of all vertices and edges appearing on any North-East path between the origin and a terminal vertex.

A *ladder diagram* is a subgraph of Γ_λ such that

(1) the origin is connected to every terminal vertex by some North-East path.
(2) every edge in the graph is on a North-East path from the origin to some terminal vertex.

An example of the grid Γ_λ and some of its ladder diagrams are shown in Figure 3. All fixed coordinates are shaded. The terminal vertices lie along the *main diagonal* of Γ_λ. As standalone objects, the ladder diagrams of Γ_λ form a poset ordered by inclusion: given two ladder diagrams $\mathcal{L}_1, \mathcal{L}_2$, we have $\mathcal{L}_1 \le \mathcal{L}_2$ if $\mathcal{L}_1$ is a subgraph of $\mathcal{L}_2$.

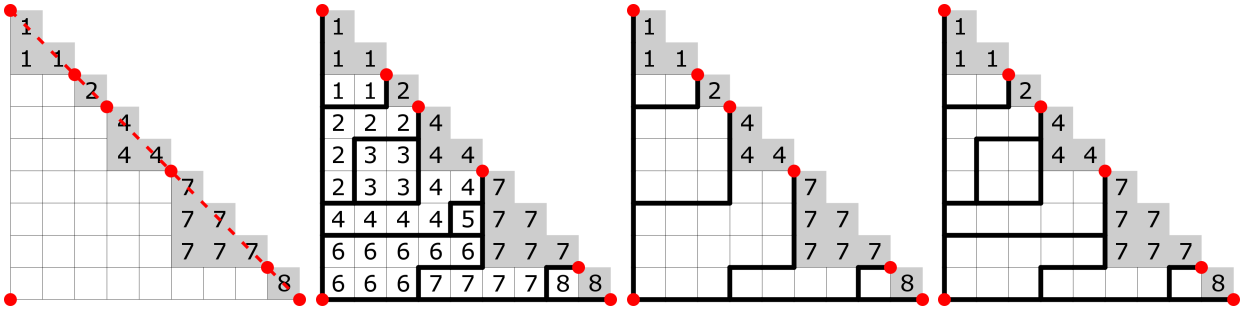

FIGURE 3. Let $\lambda = (1^2, 2^1, 4^2, 7^3, 8^1)$. From left to right: Γ_λ with origin and terminal vertices as dots and a dashed line indicating the main diagonal, ladder diagram for a point in GT_λ, ladder diagram for a 0-dimensional face (vertex), and ladder diagram for a 2-dimensional face.

Definition 2.3 (Face Lattice of Ladder Diagrams)**.** Let $\mathcal{F}(\Gamma_\lambda)$ denote the set of all ladder diagrams of Γ_λ ordered by inclusion. This may also be called the face lattice of Γ_λ.

Theorem 2.4. [ACK16, Theorem 1.9] $\mathcal{F}(GT_\lambda) \cong \mathcal{F}(\Gamma_\lambda)$.

Proof. An isomorphism is given by taking a point in GT_λ and drawing lines around adjacent groups of $x_{i,j}$ with equal value will produce a face of Γ_a. For more details, see [ACK16]. We also mention that [McA06] proves an analogous relation between $\mathcal{F}(\mathrm{GT}_\lambda)$ and the poset of GT tilings, which is essentially equivalent to $\mathcal{F}(\Gamma_\lambda)$. □

Note that $\mathcal{F}(\Gamma_\lambda)$ is graded by the number of bounded regions where k-dimensional faces correspond to ladder diagrams with k bounded regions. In fact, given a point $x \in \mathrm{GT}_\lambda$, we can determine the dimension of the minimal face containing x by mapping x to its corresponding ladder diagram. Then the number of bounded regions will be the dimension of this minimal face.

Remark 2.5. Note that the poset $\mathcal{F}(\Gamma_\lambda)$ only depends on the multiplicities a_i and not on the values of λ_i. So when examining the purely combinatorial properties of GT_λ, it suffices to consider partitions of the form $\lambda = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m})$.

3. Combinatorial Diameter

In this section, we present an exact formula for the diameter of the 1-skeleton of GT_λ which is denoted by $\mathrm{diam}(\mathrm{GT}_\lambda)$. As explained in Remark 2.5, it suffices to consider $\lambda = (1^{a_1}, \ldots, m^{a_m})$ where $a_1, \ldots, a_m \in \mathbb{Z}_{>0}$. Our proofs will indirectly work with vertices and edges of GT_λ by working with their corresponding ladder diagrams.

In order to study the diameters of the 1-skeleton of GT_λ, we need to first understand what a vertex is in terms of ladder diagrams and under what conditions two vertices are connected.

Definition 3.1. Two paths in a ladder diagram from the origin to terminal vertices are *noncrossing* if they do not meet again after their first separation.

In particular, vertices of GT_λ have ladder diagrams consisting of $m+1$ noncrossing paths. Two vertices are connected by an edge if the union of their ladder diagrams has exactly one bounded region, as shown in Figure 5. We think of traveling from one vertex to another as moving a *subpath* of the first vertex's ladder diagram. Note that such a move can alter more than one of the $m+1$ noncrossing paths.

Lemma 3.2. *Any two vertices v and w of GT_λ are separated by at most $2m-2-\delta_{1,a_1}-\delta_{1,a_m}$ edges.*

Proof. We give an algorithm to find a path between v and w of length at most $2m-\delta_{1,a_1}-\delta_{1,a_m}$. Assume that in the ladder diagram representation, vertex v corresponds to noncrossing paths

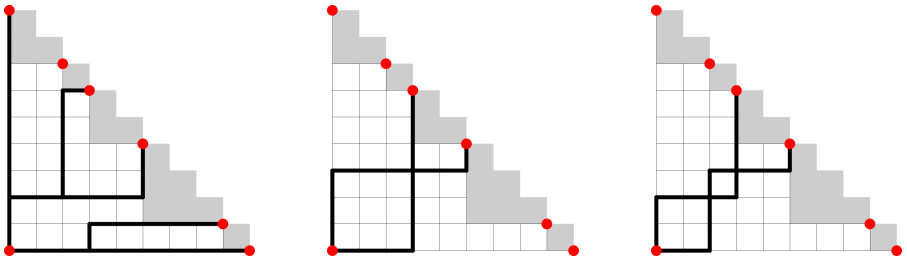

FIGURE 4. Left: 5 non-crossing paths. Middle and Right: 2 Crossing paths

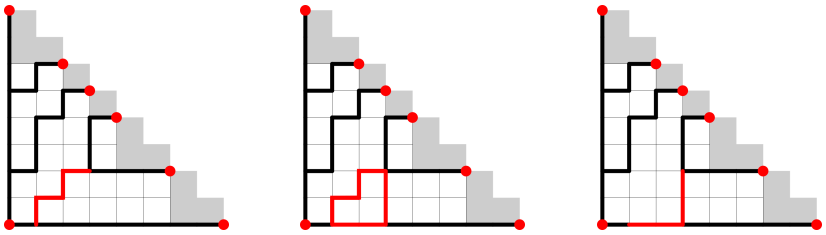

FIGURE 5. Two vertices and the edge between them.

$v_1, \ldots, v_{m-1}$ where v_j connects the origin $(0,0)$ to terminal vertex $t_j = (s_j, n - s_j)$ where $s_j = \sum_{i=0}^{j} a_i$. Similarly denote the noncrossing paths corresponding to w as $w_1, \ldots, w_{m-1}$.

Essentially, we want to change v_1 to w_1, v_2 to w_2, ..., v_{m-1} to w_{m_1}, making sure that the $m-1$ paths we have are always noncrossing, and that the common refinement before and after changing some paths has exactly one bounded region. This ensures we are always traveling along edges in GT_λ. Note that the two paths ending at t_0 and t_m are the same for every ladder diagram so we ignore them here.

Phase 1: If $a_1 = 1$, then v_1, w_1 are paths that go from $(0,0)$ to $(1, n-1)$. Therefore, there exists a unique index r_v such that path v_1 passes through both $(0, r_v)$ and $(1, r_v)$. In other words, r_v is the vertical index for v_1 to go from column 0 to column 1. Similarly we can define r_w. WLOG, assume that $r_v \geq r_w$. Because of this inequality, we know that path w_1 is contained inside of v_1 and therefore, the ladder diagram consisting of $w_1, v_1, v_2, \ldots, v_{m-1}$ has exactly 1 bounded region and it is thus an edge e of the GT polytope containing v. Let $v' = (v'_1, \ldots, v'_{m-1})$ be the other side of this edge (one side is v). Notice that the inner edge $\{(0, r_w), (1, r_w)\}$ is in the ladder diagram of e but not in the ladder diagram of v'. Therefore, $\{(0, r_w), (1, r_w)\}$ must be in the ladder diagram of v', which means $v'_1 = w_1$, since $a_1 = 1$. Similarly, we can use one move to make v_{m-1} and w_{m-1} equal if $a_m = 1$.

Phase 2: Now we describe an algorithm that takes v' to some vertex u in at most $m-1-\delta_{1,a_1}-\delta_{1,a_m}$ steps. The algorithm works as follows: for each $i = 1+\delta_{1,a_1}, \ldots, m-1-\delta_{1,a_m}$, change path v_i so that it starts at terminal vertex t_i, goes horizontally to the left until it meets and merges with path v_{i-1}. First, the ladder diagram after this change is clearly a vertex. Also, if we take the common refinement of the two ladder diagrams before and after the change, or equivalently, start with the old ladder diagram and add a new path v'_i described above, then this new path simply cuts the tile bounded by v_{i-1} and v_i into two

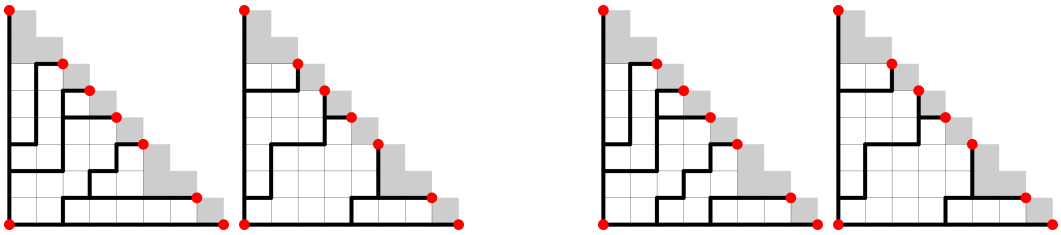

FIGURE 6. Phase 1 of the algorithm for Lemma 3.2. From left to right: the ladder diagram for v, w, v', $w'(=w)$.

parts and thus there exists an edge between these two vertices. Figure 7 shows an example of this algorithm.

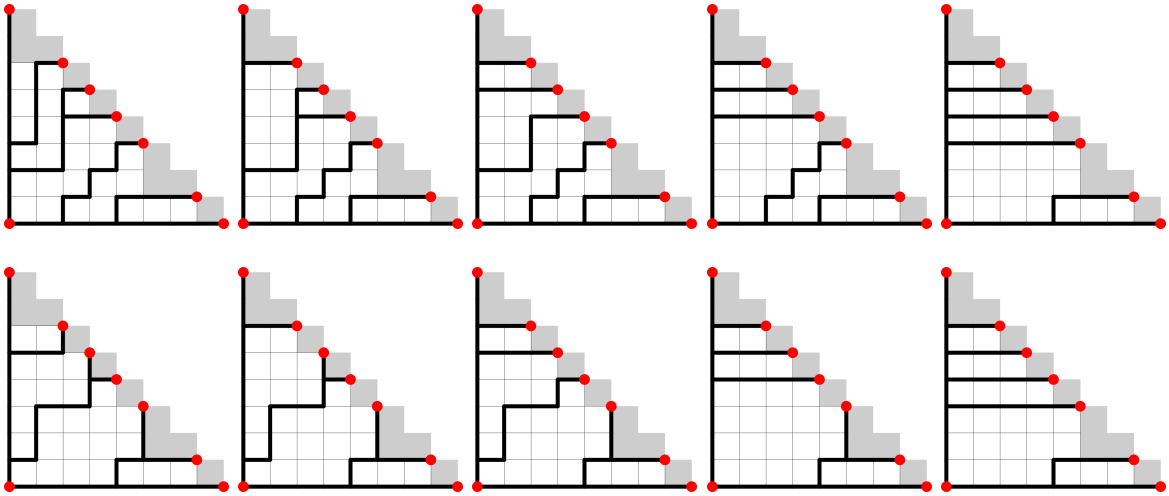

FIGURE 7. Phase 2 of the algorithm for Lemma 3.2. The first line shows steps that move from v' to u; the second line shows steps that move from w' to u.

Similarly, we can apply the same algorithm to w' to get to the same vertex u in the same number of steps. The total number of steps is at most

$$(\delta_{1,a_1} + \delta_{1,a_m}) + 2(m-1-\delta_{1,a_1} - \delta_{1,a_m}) = 2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}.$$

□

Lemma 3.3. *There exist two vertices separated by at least* $2m-2-\delta_{1,a_1}-\delta_{1,a_m}$ *edges.*

In order to prove this lemma, we need to first construct vertices in terms of ladder diagrams.

Definition 3.4 (Zigzag lattice path)**.** We construct two vertices z_h and z_v that will be used in the proof for Lemma 3.3. Let $\lambda = (1^{a_1}, \ldots, m^{a_m})$. If $a_1 > 1$, meaning that $(1, n-1)$ is not a terminal vertex, we call $(1, n-1)$ a *virtual terminal vertex*. Similarly, if $(n-1, 1)$ is not an actual terminal vertex, meaning that $a_m > 1$, we call it a virtual terminal vertex.

We will consider the ladder diagram for a vertex of GT_λ as $m-1$ southwest lattice paths from terminal vertices to the origin. For $j = 1, \ldots, m-1$, define a horizontal zigzag path, h_j, to be the path that starts at terminal vertex t_j, goes horizontally left until reaching a column

where there exists a terminal vertex or a virtual terminal vertex on it, then goes vertically down until reaching a row where there exists a terminal vertex or a virtual terminal vertex, and so on and so forth until the path reaches column 0 or row 0. Similarly define a vertical zigzag path, v_j, with the only difference that it will start vertically instead of horizontally. Finally, let z_h be the vertex of GT_λ represented by the ladder diagram $(h_1, \ldots, h_{m-1})$ and let z_v be the vertex of GT_λ represented by $(v_1, \ldots, v_{m-1})$. Figure 8 shows the construction of an example.

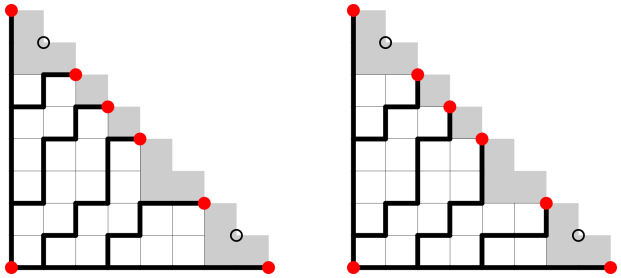

FIGURE 8. Vertices z_h (left) and z_v (right) of GT_λ with $\lambda = (1^2, 2^1, 3^1, 4^2, 5^2)$. Virtual terminal vertices are labeled as open circles.

Proof of Lemma 3.3. We will first consider the case where $a_1, a_m \geq 2$ so that the idea of the proof can be shown clearly. Afterward, we will deal with the details coming from either of them being 1.

Consider the vertices given in Definition 3.4. Assume that there is a sequence of vertices $z_h = y_0, y_1, \ldots, y_\ell = z_v$ in GT_λ such that y_k and y_{k+1} are connected by an edge. Since we can uniquely represent each vertex in this sequence as a union of $m-1$ lattice paths from the original to each terminal vertex, each step $y_k \to y_{k+1}$ can be thought of as simply changing these paths, such that the union of y_k and y_{k+1} has at most 1 bounded region. For the sake of convenience, for each $i = 1, \ldots, m-1$, let p_i be the path that goes from the origin to the terminal vertex i generically for all y_k's. For $k = 1, \ldots, \ell$, define $X_k = \{i : p_i \text{ changes as we go from } y_{k-1} \text{ to } y_k\}$. Since in each step we can have at most 1 bounded region between y_{k-1} and y_k, it is clear that X_k is of the form $\{i, i+1, \ldots, j\}$ for some $i \leq j$.

We will now show that $X_1, \ldots, X_\ell$ must satisfy the following conditions:

(1) For each $s \in [m-1]$, s appears in at least two of the sets X_i's. In other words, for each $s \in [m-1]$, there exists $1 \leq i < j \leq \ell$ such that $s \in X_i$, $s \in X_j$.
(2) For each $s \neq s' \in [m-1]$, the last time s appears in any of the sets (which is the unique index b_s such that $s \in X_{b_s}$ and $k \notin X_i$ for all $i > b_k$) is different from the last time that k' appears.
(3) If $X_k = \{i, i+1, \ldots, j\}$, then at least $j-i$ of $i, i+1, \ldots, j$ must appear in some $X_{k'}$ for $k' < k$
(4) If $X_k = \{i, i+1, \ldots, j\}$ and it is the last time that s appears, then each one of $\{i, i+1, \ldots, j\} \backslash \{s\}$ must appear in (possibly different) $X_{k'}$ for $k' < k$.

Notice that if the last time s appears is in some X_k, then p_s has already become v_s (Definition 3.4) in y_k, and stays the same for the rest of the steps. We will then explain these conditions in detail one by one.

Condition (1). If for some $s \in [m-1]$, it appears in only one set X_k, then it means when we go from vertex y_{k-1} to y_k, the path p_s is changed from v_s to h_s in exactly one step. But by construction, superimposing h_s and v_s will create at least 2 bounded regions, instead of one. Therefore, y_{k-1} and y_k are not connected by an edge, a contradiction. Therefore, each $s \in [m-1]$ appears in at least two sets.

Condition (2). If X_k is the last time that both s and s' appears, then it means that from y_{k-1} to y_k, paths $p_{s'}$ and p_s are changed to $v_{s'}$ and v_s simultaneously. However, $v_{s'}$ and v_s do not have any intersection in the interior of $\mathcal{I}_n$. So if we want to go back from y_k to y_{k-1}, we have to change v'_s and v_s simultaneously, which will create two bounded regions when we superimpose y_{k-1} and y_k, which is a contradiction.

Condition (3). Since initially, h_i and h_j do not have any intersection in the interior of $\mathcal{I}_n$, in order to change multiple paths at the same time, we need to merge these paths first. Specifically, if we want to change paths $i, i+1, \ldots, j$ simultaneously, we need to modify at least all but one of these paths to join them together.

Condition (4). This condition is crucial and it justifies our choices for z_h and z_v. Assume that $X_k = \{i, i+1, \ldots, j\}$ is the last time that s appears, meaning that when we go from vertex y_{k-1} to y_k, we change p_s to v_s. If there exists $s' \neq s \in X_k$, such that path s' hasn't changed before, we know $p_{s'} = h_{s'}$. Notice that since we change paths $i, i+1, \ldots, j$ simultaneously, all of these paths before and after the change will have an interior vertex in $\mathcal{I}_n$ in common. Therefore, v_s and $h_{s'}$ must have a common interior vertex. Since $s \neq s'$, we must have $s' = s+1$. As we have assumed $a_1, a_m \geq 2$, superimposing the ladder diagram of y_{k-1} and y_k will create at least 2 bounded regions, because of the definition of v_s and h_{s+1}. Therefore, we have a contradiction and thus all $s' \neq s$ must have already appeared at least once.

As we have proved that $X_1, \ldots, X_\ell$ satisfy all these conditions, we will jump out of the general setting of GT polytopes and look at any sequence of sets $X_1, \ldots, X_\ell$ that satisfies all these four conditions. We will show that any such sequence will have length $\ell \geq 2m-2$.

To do this, for $i = \ell, \ldots, 1$, we look at the first set X_i that is not a singleton. Say that it is X_k. If X_k is the last time that some $s \in [m-1]$ appears, then we claim that changing X_k to $\{s\}$ will still satisfy all four conditions. According to Condition (4), for each $s' \neq s$ that is in X_k, s' must have appeared before. According to Condition (2), for each $s' \neq s$, this is not the last time that s' appears so s' will appear sometime later. Therefore, condition (1) still holds after changing X_k to $\{s\}$. Condition (2) also holds because this change does not modify the indices of the sets where each $s' \in [m-1]$ appears last. Condition (3) and (4) hold trivially because we have less non-singleton sets to worry about. Another case is that X_k is not the last time that any of the paths appears last. According to condition (3), there exists $s \in X_k$ such that each one of $X_k \backslash \{s\}$ has appeared before. Similarly, we claim that all these four conditions will hold after changing X_k to $\{s\}$. For each one of $s' \in X_k \backslash \{s\}$, as

it appears before and X_k is not the last time that it appears, we know that s' will appear at least twice even after this change. The number of appearances of s does not change. Therefore, condition (1) is satisfied. Condition (2) holds because each one of $s' \in X_k$ will appear sometime later. Condition (3) and (4) hold trivially because similarly we have less non-singleton sets to worry about.

Continuing this procedure inductively, we will eventually end up with a sequence of sets $Y_1, \ldots, Y_\ell$ where each one is a singleton. As Condition (1) still holds, we conclude that $\ell \geq 2m-2$ as desired.

Now we consider the cases where a_1 and a_m may be 1. We claim that if $a_1 = 1$, then $\{1\}$ must appear as a singleton as one of the terms in $X_1, \ldots, X_\ell$ and if $a_m = 1$, then $\{m-1\}$ must appear as a singleton as one of the terms in the sequence. Notice the reasons are slightly different. If $a_1 = 1$, path h_1 has only 1 interior edge, namely $(0, n-1)-(1, n-1)$. In order for it to merge with other paths, it must change on its own first, meaning that $\{1\}$ will appear, because h_1 and $p_2, \ldots, p_{m-1}$ will never have any interior intersection. If it never merges with other paths, then it must appear as $\{1\}$ at some point so that p_1 can be actually changed from h_1 to v_1. If $a_m = 1$, at some point, path p_{m-1} must be changed to v_{m-1}, which has only 1 interior edge. This change is recorded as $\{m-1\}$ because $p_1, \ldots, p_{m-2}$ cannot have any interior intersection with v_{m-1}.

Therefore, if $a_1 = 1$ (or/and $a_m = 1$), in the sequence $X_1, \ldots, X_\ell$, we can take out the singleton $\{1\}$ (or/and $\{m-1\}$) and delete all other 1's (or/and $m-1$'s) in the sets. The remaining sequence will satisfy the four conditions mentioned above. We have taken out at least $\delta_{1,a_1} + \delta_{1,a_m}$ singletons of $\{1\}$ and $\{m-1\}$ and there are $m - 1 - \delta_{1,a_1} - \delta_{1,a_m}$ paths remaining. So by the same argument, the total length of the sequence

$$\ell \geq (\delta_{1,a_1} + \delta_{1,a_m}) + 2(m - 1 - \delta_{1,a_1} - \delta_{1,a_m}) = 2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}$$

as desired. □

Lemma 3.2 and Lemma 3.3 together prove our first result, Theorem 1.1.

Theorem 1.1 (Diameter of 1-skeleton)**.** *diam*$(GT_\lambda) = 2m - 2 - \delta_{1,a_1} - \delta_{1,a_m}$.

4. Combinatorial Automorphisms

In this section, we completely describe the combinatorial automorphism group of GT_λ which is denoted by $\mathrm{Aut}(\mathrm{GT}_\lambda)$. A combinatorial automorphism of a polytope is a permutation of its faces that preserves inclusion. Such a permutation can be thought of as an automorphism of the face lattice. For the sake of brevity, we will refer to combinatorial automorphisms simply as automorphisms. In Section 4.1 we present the generators of the automorphism group and the relations between them. In Section 4.2 we discuss *facet chains* which are then used in Section 4.3 to show that our generators form the entire automorphism group. As explained in Remark 2.5, it suffices to consider $\lambda = (1^{a_1}, \ldots, m^{a_m})$ where $a_1, \ldots, a_m \in \mathbb{Z}_{>0}$.

4.1. **Automorphisms.** We begin by identifying the generators of $\mathrm{Aut}(\mathrm{GT}_\lambda)$. By Theorem 2.4, to show that a map is an automorphism of GT_λ, it suffices to show that the map is an automorphism of $\mathcal{F}(\Gamma_\lambda)$.

Proposition 4.1 (The Corner Symmetry)**.** *For any λ, there is an order 2 automorphism μ of $\mathcal{F}(\Gamma_\lambda)$ which acts by exchanging two pairs of edges near $(0,0)$: if a ladder diagram contains $\{(0,0),(1,0)\}$ and $\{(1,0),(1,1)\}$, then μ acts on this diagram by replacing these edges with $\{(0,0),(0,1)\}$ and $\{(0,1),(1,1)\}$ and vice versa.*

Proof. Let $\mu : \mathbb{R}^{\frac{n(n+1)}{2}} \to \mathbb{R}^{\frac{n(n+1)}{2}}$ be the linear map that sends $x_{n,1} \mapsto x_{n,2} + x_{n-1,1} - x_{n,1}$ and acts as the identity on all other $x_{i,j}$. Since $x_{n-1,1} \leq x_{n-1,1} + x_{n,2} - x_{n,1} \leq x_{n,2}$, all necessary inequalities are satisfied so $\mu(\mathrm{GT}_\lambda) \subset \mathrm{GT}_\lambda$. Since $\mu^2 = Id$, the previous inclusion is actually an equality. Combining this fact with the fact that μ is an affine transformation, we see μ induces a combinatorial automorphism on GT_λ, which we shall abuse notation and call μ. Examining what μ does to equalities shows it acts in the described manner on ladder diagrams.

□

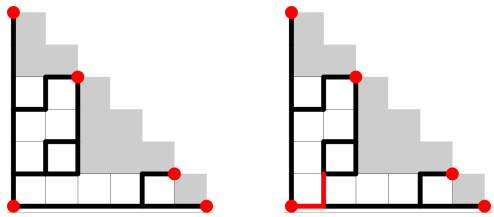

FIGURE 9. Action of μ on a ladder diagram.

The next map is very similar to the Corner Symmetry μ except it occurs at the kth terminal vertex instead of at the origin. The argument is very similar to Prop 4.1, so we will omit some details.

Proposition 4.2 (The k-Corner Symmetry)**.** *Let the kth terminal vertex be $t_k = (n-i,i)$ and suppose that $a_k, a_{k+1} \geq 2$. Then there is an order 2 automorphism μ_k which acts by exchanging two pairs of edges near t_k: if a ladder diagram contains $\{(n-i,i),(n-i,i-1)\}$ and $\{(n-i,i-1),(n-i-1,i-1)\}$, then μ_k acts on this diagram by replacing these edges with $\{(n-i,i),(i-1,i)\}$ and $\{(n-i-1,i),(n-i-1,i-1)\}$.*

Proof. Let $x_{i',j'}$ be the coordinate in Figure 1 which is immediately adjacent to the values λ_k and λ_{k+1}. Let $\mu_k : \mathbb{R}^{\frac{n(n+1)}{2}} \to \mathbb{R}^{\frac{n(n+1)}{2}}$ be the linear map that sends $x_{i',j'} \mapsto x_{i',j'-1} + x_{i'+1,j'} - x_{i',j'}$ and acts as the identity on all other $x_{i,j}$. Since $a_k, a_{k+1} \geq 2$, we have

$$\lambda_k \leq x_{i',j'-1} \leq x_{i',j'-1} + x_{i'+1,j'} - x_{i',j'} \leq x_{i'+1,j'} \leq \lambda_{k+1}$$

so all 4 constraints on the coordinate $x_{i',j'}$ are satisfied. Since $\tau^2 = 1$, the same argument as in Prop. 4.1 gives an automorphism τ with the desired properties.

□

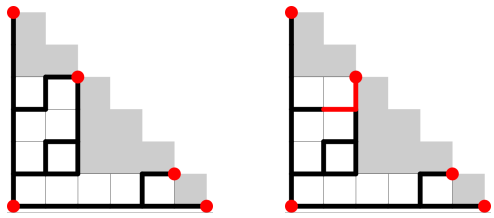

FIGURE 10. Action of μ_k on a ladder diagram.

Proposition 4.3 (Symmetric Group Symmetry)**.** *If $a_1 = 1$, then $S_{a_2} \subset Aut(GT_\lambda)$. Similarly if $a_m = 1$, then $S_{a_{m-1}} \subset Aut(GT_\lambda)$.*

Proof. Number the topmost a_2 horizontal edges in the first column of Γ_λ from 1 through a_2. Let σ be any element in S_{a_2}. Then σ acts on each ladder diagram by permuting the horizontal edges in the first column. Specifically, for a given ladder diagram, if there is a path passing through the ith horizontal edge and $\sigma(i) = j$, then σ acts on this ladder diagram by changing this path to go through the jth horizontal edge. Note that there is only one way to change the vertical edges of the modified path to yield a valid ladder diagram. Again, σ's action on $\mathcal{F}(\mathrm{GT}_\lambda)$ has finite order so σ is a bijection on the underlying set of $\mathcal{F}(\mathrm{GT}_\lambda)$. Furthermore, σ preserves inclusions of ladder diagrams so σ is an automorphism of $\mathcal{F}(\mathrm{GT}_\lambda)$. □

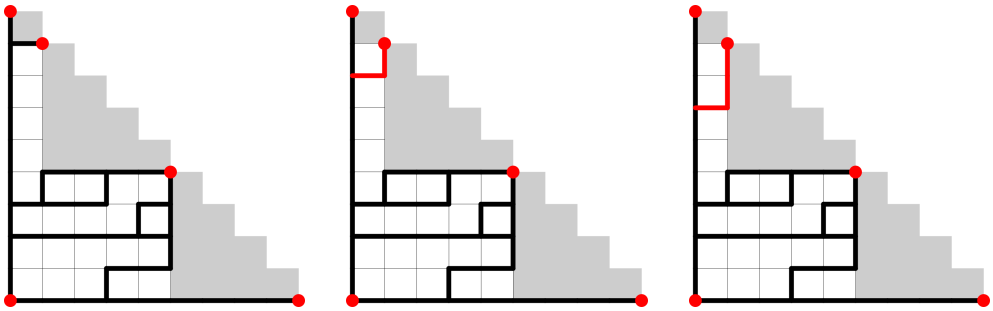

FIGURE 11. Action of $(123) \in S_4$ on a ladder diagram.

Proposition 4.4 (The Flip Symmetry)**.** *Suppose that $\lambda = \lambda'$. Then there is an order 2 automorphism ρ of $\mathcal{F}(GT_\lambda)$ which reflects ladder diagrams over the line $y = x$.*

Proof. Recall that if $\lambda = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m})$, then $\lambda' := (1^{a_m}, 2^{a_{m-1}}, \ldots, m^{a_1})$. So if $\lambda = \lambda'$, we have $a_i = a_{m-i+1}$ for $1 \le i \le m$. In other words, the terminal vertices of Γ_λ are symmetric about the line $y = x$. Thus ρ sends ladder diagrams to valid ladder diagrams. Since $\rho^2 = 1$ and ρ preserves inclusions of ladder diagrams, ρ is an automorphism of $\mathcal{F}(\Gamma_\lambda)$. □

Remark. When $\lambda = (1^{a_1}, 2^{a_2}, \ldots, m^{a_m})$, this symmetry is actually affine: it can be realized as the map $f(x) = -Px + m \cdot \mathbf{1}$, where P is a permutation matrix, and $\mathbf{1}$ is the vector with each entry equal to 1. However, for λ of a different form, this is no longer true, even in small cases. For example when $\lambda = (1, 2, 4)$, a straightforward computation shows that ρ is no longer affine.

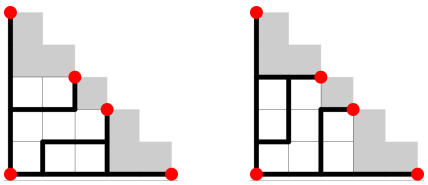

FIGURE 12. Action of ρ on ladder diagram.

Proposition 4.5 (The $m = 2$ Rotation Symmetry)**.** *Suppose that $m = 2$. There is an order 2 automorphism τ on $\mathcal{F}(GT_\lambda)$ that rotates all paths from $(0,0)$ to t_1 by $180°$ into paths from t_1 to $(0,0)$.*

Proof. The map τ takes ladder diagrams to valid ladder diagrams. Since $\tau^2 = 1$ and τ preserves inclusions of ladder diagrams, τ is an automorphism of $\mathcal{F}(\mathrm{GT}_\lambda)$. □

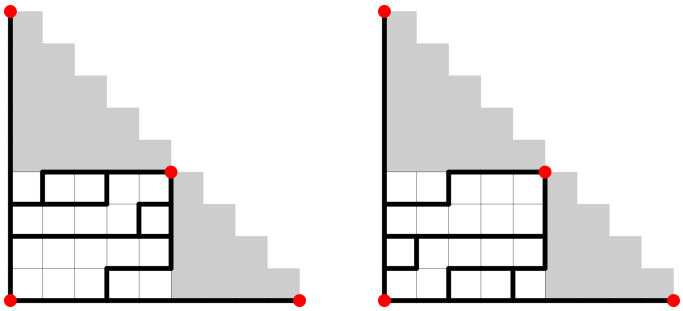

FIGURE 13. Action of τ on a ladder diagram.

Proposition 4.6 (The $m = 2$ Vertex Symmetry)**.** *When $m = 2$, there are two paths p_1 and p_2 to the terminal vertex t_1 that turn exactly once. The map α sending these two paths to each other is an order 2 automorphism of GT_λ.*

Proof. Similar to the proofs above, $\alpha^2 = 1$ and it clearly preserves inclusion. Therefore, it is an automorphism of $\mathcal{F}(\mathrm{GT}_\lambda)$. □

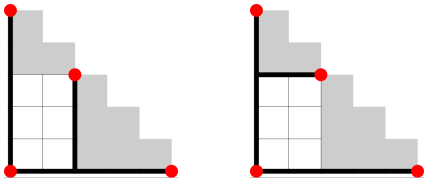

FIGURE 14. Paths p_1 and p_2.

The ladder diagram formed by p_1 is a vertex of GT_λ that is connected to every other vertex by an edge, and similarly for p_2. The map α simply exchanges these two vertices.

It is clear that the group formed by these possible generators is contained in $\mathrm{Aut}(\mathrm{GT}_\lambda)$. Note that if $m = 1$, then the polytope is a single point with only the trivial automorphism. If

$m = 2$ and either $\lambda = (1, 2^{a_2})$ or $\lambda = (1^{a_1}, 2)$, then the polytope is a simplex and its automorphism group is the symmetric group. In all other cases, either $m = 2$ and $a_1, a_2 \geq 2$, or $m \geq 3$. For both of these cases, we will describe the group generated by the automorphisms listed above and finish proving that this group is the entire automorphism group in Section 4.3.

Theorem 1.2 (m=2 Automorphisms)**.** *Suppose $\lambda = (1^{a_1}, 2^{a_2})$ and $a_1, a_2 \geq 2$. If $a_1 = a_2 = 2$, then*

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2.$$

Otherwise,

$$Aut(GT_\lambda) \cong D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}}.$$

Proof of inclusion. We show that if $a_1 = a_2 = 2$, then $D_4 \times \mathbb{Z}_2 \subseteq \mathrm{Aut}(\mathrm{GT}_\lambda)$ and otherwise $D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}} \subseteq \mathrm{Aut}(\mathrm{GT}_\lambda)$.

Suppose $a_1 \neq a_2$. Collecting the generators applicable in this case from previous propositions, we have the subgroup of automorphisms $\langle \mu, \mu_1, \tau, \alpha \rangle$, and note these generators satisfy the following relations $\mu^2 = \mu_1^2 = \tau^2 = \alpha^2 = 1, \mu\tau = \tau\mu_1$ and all other variables commute. The subgroup $\langle \tau\mu, \mu \rangle$ is D_4. The generator α commutes with all other generators, so the resulting group is isomorphic to $D_4 \times \mathbb{Z}_2$.

Assuming $a_1 = a_2 \geq 3$, we have the subgroup described in the previous case $a_1 \neq a_2$, but with the additional generator ρ. Note ρ commutes with all of these generators from the previous case, so the resulting group is $D_4 \times \mathbb{Z}_2^2$. When $a_1 = a_2 = 2$, note that $\rho = \tau$, resulting in a smaller subgroup. □

We take care in specifying how to write down the composition of elements of $\mathrm{Aut}(\mathrm{GT}_\lambda)$ as a tuple. This is not as straightforward as in the $m = 2$ case because the Flip Symmetry does not act locally and does not commute with the other symmetries.

In the statement of Theorem 1.3, elements of $\mathrm{Aut}(\mathrm{GT}_\lambda)$ are tuples $(\sigma_1, \sigma_2, z_1, \ldots, z_{r_1}, z_{r_1+1}, z_{r_1+2})$. Here $\sigma_1 \in S_{a_2}^{\delta_{1,a_1}}$ and $\sigma_2 \in S_{a_{m-1}}^{\delta_{1,a_m}}$ correspond to the Symmetric Group Symmetry and $z_1, \ldots, z_{r_1} \in \mathbb{Z}_2$ correspond to the k-Corner Symmetries. Finally $z_{r_1+1} \in \mathbb{Z}_2$ corresponds to the Corner Symmetry, $z_{r_1+2} \in \mathbb{Z}_2$ corresponds to the Flip Symmetry. Let $g \in S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}$ be such that $(\sigma_1, \sigma_2, z_1, \ldots, z_{r_1}, z_{r_1+1}, z_{r_1+2}) =: (g, z')$. Then $(g, z_{r_1+2}) \cdot (g', z'_{r_1+2}) = (g\varphi(z_{r_1+2})(g'), z_{r_1+2}z'_{r_1+2})$ where $\varphi(0)$ is the identity map on $S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}}$ and $\varphi(1)$ is the map sending $(\sigma_1, \sigma_2, z_1, \ldots, z_{r_1}, z_{r_1+1}) \mapsto (\sigma_2, \sigma_1, z_{r_1}, \ldots, z_1, z_{r_1+1})$. This is formally stated in Theorem 1.3.

Theorem 1.3 ($m \geq 3$ Automorphisms)**.** *Suppose $\lambda = 1^{a_1} \ldots m^{a_m}$ and $m \geq 3$. Let r_1 be the number of k such that $a_k, a_{k+1} \geq 2$. Let $r_2 = 1$ if $\lambda = \lambda'$ and let $r_2 = 0$ otherwise. Then*

$$Aut(GT_\lambda) \cong (S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}) \ltimes_\varphi \mathbb{Z}_2^{r_2}$$

where if $r_2 = 1$, then $\varphi : \mathbb{Z}_2 \to Aut(S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1})$ sends the nonidentity element of $\mathbb{Z}_2$ to the map sending $(\sigma_1, \sigma_2, z_1, \ldots, z_{r_1}, z_{r_1+1}) \mapsto (\sigma_2, \sigma_1, z_{r_1}, \ldots, z_1, z_{r_1+1})$.

Proof of inclusion. We show that $(S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}) \ltimes_\varphi \mathbb{Z}_2^{r_2} \subseteq \mathrm{Aut}(\mathrm{GT}_\lambda)$.

By the previous propositions, our generators are $\mu, \mu_1, \ldots, \mu_{m-1}, S_{a_2}, S_{a_{m-1}}$ with S_{a_2} or $S_{a_{m-1}}$ possibly omitted. Whichever symmetries are present commute since they act on disjoint sets of edges in the ladder diagrams.

If λ is reverse symmetric, we also have the generator ρ. Note that since ρ flips every ladder diagram about $y = x$, ρ satisfies the following commutation relations: $\rho\mu = \mu\rho$, $\rho\mu_i = \mu_{m-i}\rho$, and for $\sigma \in S_{a_2}, S_{a_{m-1}}$ (a_2 and a_{m-1} must be equal for λ to be reverse symmetric),$\rho\sigma = \sigma\rho$. These relations are enough to give a subgroup of the stated form. □

In the following sections, we finish the proofs of Theorems 1.2 and 1.3 by showing the group formed by our generators is the entire automorphism group by bounding the size of $\mathrm{Aut}(\mathrm{GT}_\lambda)$ by the size of the group formed by our generators. We will examine the action of any combinatorial automorphism on $\mathcal{F}(\Gamma_\lambda)$ and apply the Orbit-Stabilizer theorem. In order to do so, we first develop ways to classify and partition the facets of GT_λ.

4.2. **Chains of Facets.** In this section, we present a useful model for the facets of GT_λ, define the *facet chains* of Γ_λ, and discuss how an automorphism can act on facet chains. This line of reasoning is motivated by the following observation.

Lemma 4.7. *An automorphism of GT_λ is determined by where it sends the facets of GT_λ or equivalently, where it sends the ladder diagrams of facets.*

Proof. This follows from the fact that for a general polytope P, every face of P can be written as an intersection of the facets of P. Thus specifying the image of each facet suffices to specify the image of any face. □

Definition 4.8. We define the *interior edges* of Γ_λ to be all edges of the form $\{(s_j, n - s_{j+1}), (s_j, n - s_{j+1} + 1)\}$ or $\{(s_j, n - s_{j+1}), (s_j + 1, n - s_{j+1})\}$ and all edges lying inside Γ_λ. All other edges of Γ_λ are considered *boundary edges*.

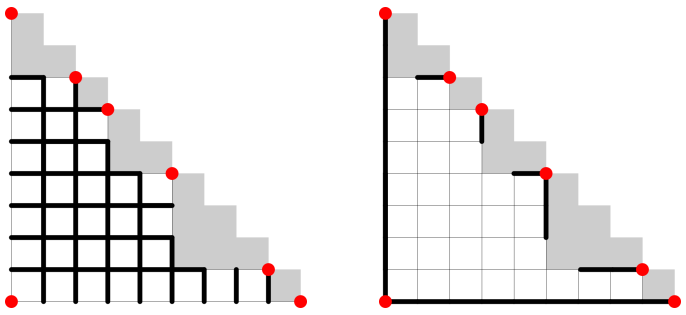

FIGURE 15. Left: interior edges of Γ_λ. Right: boundary edges of Γ_λ.

Proposition 4.9. *The facets of GT_λ are in bijection with the interior edges of Γ_λ.*

Proof. Since ladder diagrams in $\mathcal{F}(\Gamma_\lambda)$ are graded by number of bounded regions, any facet will correspond to a ladder diagram with all possible edges except one which will be an interior edge. The bijection is given by mapping a facet F to the single interior edge not contained in its associated ladder diagram. For the other direction, note that any ladder diagram missing a single interior edge is a valid ladder diagram and represents a facet. □

We will often represent a facet $F \in \mathrm{GT}_\lambda$ by its corresponding interior edge which we denote by $e(F)$. A face is not contained in facet F iff its ladder diagram contains edge $e(F)$.

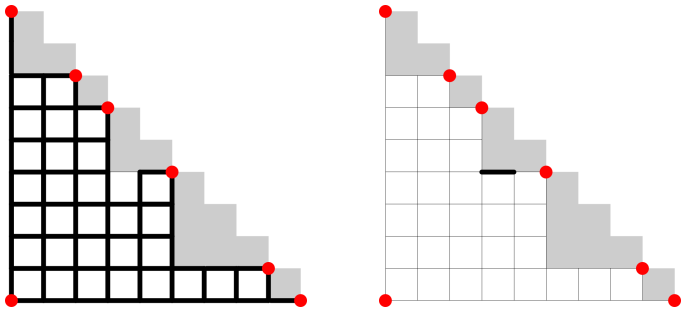

FIGURE 16. Left: ladder diagram of a facet. Right: complement of ladder diagram.

Definition 4.10. Two facets are called *dependent* if their intersection is a $d-3$ dimensional face.

Given two facets F_1 and F_2, a necessary condition for F_1 and F_2 to be dependent is for them to be arranged as shown in Figure 17. As shown in Figure 18, this is not always sufficient if there are already forced equalities amongst the 4 shaded coordinates.

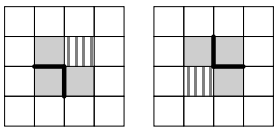

FIGURE 17. The gray boxes indicate coordinates $x_{i,j}$ that are equal on each facet. The dashed box indicates the coordinate forced to be equal to the other three.

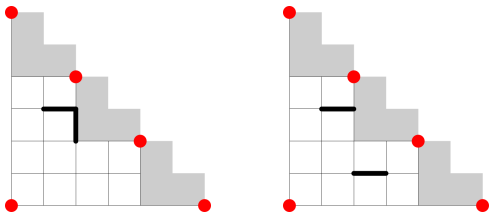

FIGURE 18. Examples of facets that are not dependent.

Facet dependencies naturally partitions the facets of GT_λ into *facet chains* which are easy to represent visually.

Definition 4.11 (Facet Chains)**.** A *facet chain* $C = (F_1, \ldots, F_l)$ of length l is an ordered list of facets $F_1, \ldots, F_l$ such that F_j is dependent on F_{j-1} and facets F_1, F_l are not dependent on any facets not in C. Visually, a chain is a set of edges $e(F_j)$ of Γ_λ forming a zig-zag pattern. The facets $F_1, \ldots, F_l$ are ordered such that $e(F_1)$ has the smallest x-coordinate.

Let $\mathcal{C}$ denote the set of facet chains of Γ_λ. These chains partition the interior edges of Γ_λ. After we have the partition, it is then natural to study the relations between these facet chains. We define the notion of adjacency between facet chains as follows.

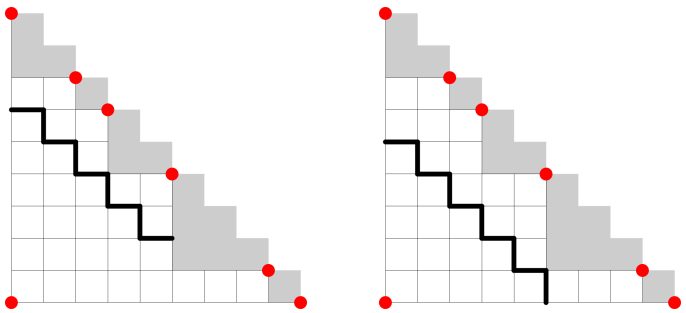

FIGURE 19. Two adjacent facet chains.

Definition 4.12 (Adjacent Chains)**.** Two different chains $C, C' \in \mathcal{C}$ are *adjacent* if there exists $F_1, F_2 \in C$ and $J_1, J_2 \in C'$ such that $F_1 \bigcap F_2 = J_1 \bigcap J_2$. We say C and C' are adjacent at k points if there are k distinct sets of facets F_1, F_2, J_1, J_2 with $F_1 \bigcap F_2 = J_1 \bigcap J_2$.

Visually, two chains are adjacent iff one chain sits directly to the North-East of the other chain as shown in Figure 19.

We are now ready to study how each automorphism $\phi \in \mathrm{Aut}(\mathrm{GT}_\lambda)$ acts on facet chains. Specifically, the action of an automorphism $\phi \in \mathrm{Aut}(\mathrm{GT}_\lambda)$ can be extended to sets of facets. For any sets of facets X_1 and X_2, we say $\phi(X_1) = X_2$ if the restriction of ϕ to X_1 is a bijection between sets X_1 and X_2. In particular, we will often abuse notation and write $\phi(C_1) = C_2$, thinking of chains C_1 and C_2 as sets of facets. We now state a few simple lemmas that will be essential to our proofs of Theorems 1.2 and 1.3.

Lemma 4.13. *Let C be a facet chain and $\phi \in Aut(GT_\lambda)$. Then $\phi(C)$ is a facet chain of the same length as C.*

Lemma 4.14. *Let C and C' be adjacent facet chains and $\phi \in Aut(GT_\lambda)$. Then $\phi(C)$ and $\phi(C')$ are also adjacent facet chains.*

Proof of Lemma 4.13 and 4.14. Dependency between facets and adjacency between facet chains can be realized as properties of coatoms in the face lattice of the GT polytope. Thus, the lemmas followed directly from the definition of an automorphism. □

Given a notion of adjacency between facet chains, it is natural to examine the adjacency graph and how an automorphism can act on it.

Definition 4.15. (Adjacency Graph) For any λ, let $\mathcal{G}_\lambda$ be the adjacency graph of the chains of Γ_λ with length ≥ 2. More precisely, the nodes of $\mathcal{G}_\lambda$ are the chains in $\mathcal{C}$ of length ≥ 2 and there is an edge between two nodes iff their corresponding chains are adjacent. Note that $\mathcal{G}_\lambda$ is connected. We make $\mathcal{G}_\lambda$ a rooted planar tree by letting its root be C_0 and giving its nodes the ordering from their corresponding chains in Γ_λ.

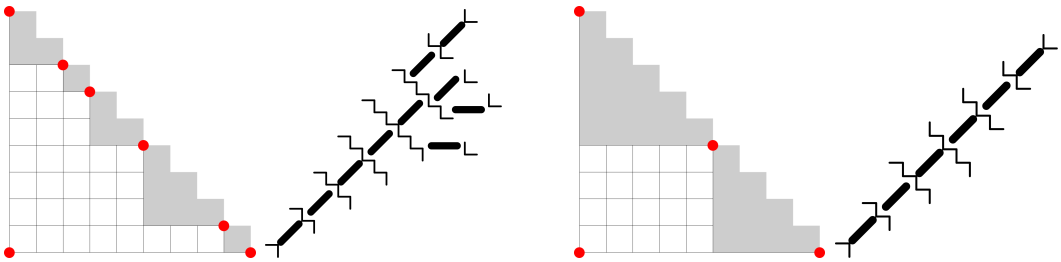

FIGURE 20. Left: $\lambda = (1^2, 2^1, 3^2, 4^3, 5^1)$ and its adjacency graph. Right: $\lambda = (1^5, 2^4)$ and its adjacency graph.

Note that the length 2 chains are exactly the chains with less than two neighbors in $\mathcal{G}$ since longer chains will have an adjacent chain above or below. Thus, the length 2 chains are the leaves of $\mathcal{G}_\lambda$.

Unless specified otherwise, when talking about the action of an automorphism on nodes of $\mathcal{G}_\lambda$, we are treating chains as sets of facets so we will not specify how the automorphism acts on the facets of chains. So if we say that an automorphism fixes a node of $\mathcal{G}_\lambda$, we are not specifying whether the automorphism flips the chain.

Lemma 4.16. *Given facet chains $C = (F_1, \ldots, F_k)$ and $C' = (F'_1, \ldots, F'_k)$ and $\phi \in Aut(GT_\lambda)$ such that $\phi(C) = C'$, either $\phi(F_i) = F'_i$ for all $1 \leq i \leq k$ or $\phi(F_i) = F'_{k+1-i}$ for all $1 \leq i \leq k$.*

Proof. If $\phi(F_1) = F'_1$, the chain of dependencies of the F_i will determine the image of each F_i. More specifically, $\phi(F_2)$ must be dependent on $\phi(F_1) = F'_1$, so we must have $\phi(F_2) = F'_2$, and so forth. Else if $\phi(F_1) \neq F'_1$, we must have $\phi(F_1) = F'_k$ since F_1 and F_k are the only facets dependent on exactly one other facet. Again the chain of dependencies implies that $\phi(F_2)$ must be dependent on $\phi(F_1) = F'_k$, so we must have $\phi(F_2) = F'_{k-1}$, and so forth. □

Under the assumptions of Lemma 4.16, if $\phi(F_i) = F'_{k+1-i}$ for all $1 \leq i \leq k$, then we say C is mapped to the *flip* of C'. In particular if $\phi(C) = C$, then ϕ either flips or does not flip C. We call this the *orientation* of chain C under ϕ.

Lemma 4.17. *Suppose C and C' are adjacent facet chains with distinct facets $F_1, F_2, F_3, F_4 \in C$ and $J_1, J_2, J_3, J_4 \in C'$ such that $F_1 \cap F_2 = J_1 \cap J_2$ and $F_3 \cap F_4 = J_3 \cap J_4$. Equivalently, C and C' are adjacent at ≥ 2 points. Given an automorphism $\phi \in Aut(GT_\lambda)$ such that $\phi(C) = C$ and $\phi(C') = C'$, C is flipped under ϕ iff C' is flipped under ϕ.*

Proof. Suppose ϕ does not flip C_i, so $\phi(F_1) = F_1$ and $\phi(F_2) = F_2$. Since $F_1 \bigcap F_2 = J_1 \bigcap J_2$, and the intersection of any pair of facets in a chain is unique, we must have $\phi(J_1), \phi(J_2) \in \{J_1, J_2\}$.

Suppose for contradiction that $\phi(J_1) = J_2$ implying C_j is flipped. Note flipping a chain can preserve at most one adjacent pair of facets in that chain, so we must have $\phi(\{J_3, J_4\}) \neq \{J_3, J_4\}$. Again the intersection of two facets in a chain is unique, implying $\phi(J_3) \bigcap \phi(J_4) \neq J_3 \bigcap J_4$, but $\phi(J_3) \bigcap \phi(J_4) = \phi(F_3) \bigcap \phi(F_4) = F_3 \bigcap F_4 = J_3 \bigcap J_4$, a contradiction. The reverse direction is similar. □

Lemma 4.18. *If an automorphism ϕ fixes every node of $\mathcal{G}_\lambda$, then all nodes corresponding to chains of length > 2 have the same orientation under ϕ.*

Proof. By Lemma 4.17, it only remains to prove that if two adjacent chains of length > 2 are adjacent at one point, then they have the same orientation under ϕ. Notice that here when we talk about a point at which two adjacent chains are adjacent, we actually mean a $d-3$ dimensional face as the intersection. See Figure 21 for the natural correspondence.

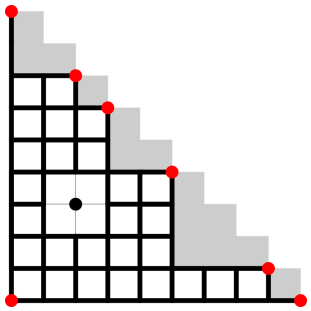

FIGURE 21. The correspondence between an interior point of Γ_λ and the ladder diagram for a $d-3$ dimensional face

Consider the set of points at which two chains with length greater than 2 can be adjacent. These are the points inside Γ_λ, excluding the corner points near terminal vertices, whose nearby parts a_i, a_{i+1} have size at least 2. We partition these points into sets as follows: two points p_1, p_2 are in the same set if they lie along the same diagonal and if all of the points lying on this diagonal and lying between p_1, p_2 are inside Γ_λ. See Figure 22 for an example.

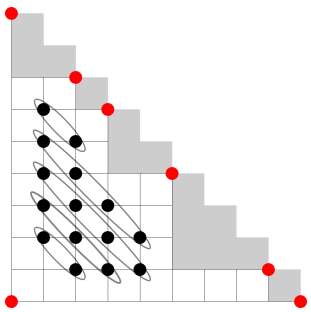

FIGURE 22. Partition of interior points

If two chains are adjacent, then the points at which they are adjacent correspond to one of these sets. So we are only concerned with adjacencies corresponding to sets of size 1. There only exist sets of size 1 corresponding to adjacencies between chains of length > 2 if either $a_1 = 2$ or $a_m = 2$. See Figure 23 for an example.

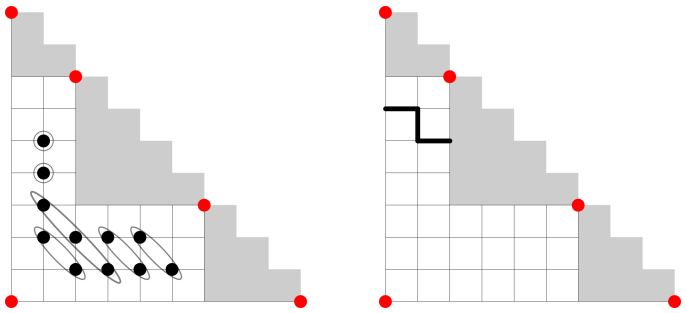

FIGURE 23. Left: partition of interior points when $a_1 = 2$. Right: example of length 3 chain D_1.

WLOG assume $a_1 = 2$. Let $D_1, \dots, D_{a_2-2}$ denote the length 3 chains such that D_1 is adjacent to the length 2 chain near t_1 and D_k is adjacent to D_{k+1} for $1 \le k \le a_2 - 3$. Once the facets in D_1 are fixed, the intersection of the two leftmost facets in D_2 is fixed since this equals the intersection of the two rightmost facets in D_1. So D_1 could not have been flipped, and D_2 is also not flipped. Similarly, we can argue that $D_3, \dots, D_{a_2-2}$ are not flipped. If D_{a_2-2} is adjacent to a chain of length > 2, repeating this argument shows that this chain cannot be flipped. In particular, all of these chains have the same orientation. □

4.3. **The Automorphism Group.** In this section, we finish the proofs of Theorems 1.2 and 1.3. Recall that in Section 4.1, we showed that the desired groups of automorphisms are contained in $\mathrm{Aut}(\mathrm{GT}_\lambda)$. We will now show equality by bounding the size of $\mathrm{Aut}(\mathrm{GT}_\lambda)$. We will begin with a few preliminaries.

For group G acting on the set of facets of GT_λ, we adopt the following conventions:

- For any set of facets X, we denote the stabilizer of X by $G(X) := \{g \in G : g \cdot x = x \ \forall x \in X\}$. We denote the orbit of X with respect to a subgroup $H \subset G$ by $O_H(X)$.
- Let $|G|$ denote the order of a group G.
- For any set of facets X, let X^c denote the set of facets not in X.

We will analyze the action of an automorphism on the facets of GT_λ and apply the Orbit-Stabilizer theorem, which states that $|G| = |O_G(X)||G(X)|$. Below we finish the proof of Theorem 1.2.

Proof of Thm. 1.2 cont. In Section 4.1, we showed that $D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}} \subseteq \mathrm{Aut}(\mathrm{GT}_\lambda)$. Now we show that the order of $\mathrm{Aut}(\mathrm{GT}_\lambda)$ is at most the order of this group.

Let $\phi \in \mathrm{Aut}(\mathrm{GT}_\lambda)$. By Lemma 4.13, ϕ preserves lengths of chains, so we may consider the action of ϕ on chains of the same length. There are two facets in length 1 chains. Let X_1 denote the set of these two facets. Let $C_0 = (F_{0,0}, F_{0,1})$ denote the length 2 chain near the origin and let $C_2 = (F_{2,0}, F_{2,1})$ denote the length 2 chain near t_1. Let $X_{2,0} := \{F_{0,0}, F_{0,1}\}$ and $X_{2,2} := X_{2,0} \cup \{F_{2,0}, F_{2,1}\}$.

To preserve lengths of chains, ϕ must send $F_{0,0}$ to a facet in a length 2 chain so $|O_G(F_{0,0})| \leq 4$. If ϕ fixes $F_{0,0}$, then ϕ fixes $F_{0,1}$ and we have $|O_{G(X_{2,0})}(F_{2,0})| \leq 2$ and $|O_{G(X_{2,0}\cup\{F_{2,0}\})}(F_{2,1})| = 1$. Any $\phi \in G(X_{2,2})$ can either fix or exchange the facets in X_0 so $|O_{G(X_{2,2})}(X_0)| \leq 2$.

Recalling Definition 4.15, if ϕ fixes all chains of length ≤ 2, then ϕ fixes every node in $\mathcal{G}_\lambda$. By Lemma 4.16, the orientation of all non-leaves must be the same. Now, we must show that $\phi \in G(X_0 \cup X_{2,2})$ can only flip all non-leaves if $\lambda = \lambda'$. Note that the corresponding tree $\mathcal{G}_\lambda$ is a path. If we have such a ϕ flipping each chain, we can show that every chain is symmetric about the line $y = x$ by inducting on the length from C_0. Every chain being symmetric about $y = x$ implies the entire diagram is symmetric about $y = x$, implying $\lambda = \lambda'$.

Therefore, $|O_{G(X_0\cup X_{2,2})}(F')| \leq 2^{\delta_{a_1,a_2}}$ and $|G(X_0 \cup X_{2,2} \cup F')| = 1$.

So we have

$$|G| = |O_G(X_{2,0})||O_{G(X_{2,0})}(X_{2,2}\backslash X_{2,0})||O_{G(X_{2,2})}(X_0)||O_{G(X_0\cup X_{2,2})}(F')||G(X_0 \cup X_{2,2} \cup \{F'\})|$$
$$\leq 4 \cdot 2 \cdot 2 \cdot 2^{\delta_{a_1,a_2}} \cdot 1.$$

We conclude $\mathrm{Aut}(\mathrm{GT}_\lambda) \cong D_4 \times \mathbb{Z}_2 \times \mathbb{Z}_2^{\delta_{a_1,a_2}}$. □

Remark. It is natural to ask how the automorphism group of GT_λ is related to the automorphism group of the 1-skeleton of GT_λ. Clearly $\mathrm{Aut}(\mathrm{GT}_\lambda)$ is contained within the automorphism group of the 1-skeleton, but we do not always have equality. If $\lambda = (3, 3)$, then $\mathrm{Aut}(\mathrm{GT}_\lambda)$ has size 16 while the automorphim group of the 1-skeleton has size 32. Numerical computations suggest that we may have equality in all other cases when $m = 2$.

The proof of Theorem 1.3 uses similar arguments. We begin by establishing some useful lemmas about the image of chains of length 1 or 2 under any automorphism.

We label the length 1 chains in Γ_λ as follows: let $\mathcal{C}_1$ be the set of length 1 chains occurring to the left of t_1 and let $\mathcal{C}_2$ be the set of length 1 chains occurring to the right of t_{m-1}. Let X_1 denote the set of facets in chains $\mathcal{C}_1 \cup \mathcal{C}_2$. Figure 24 shows all the length 1 chains visually. Let $D_1 \in \mathcal{C}_1$ and $D_2 \in \mathcal{C}_2$ be the length 1 chains with smallest y-coordinate and x-coordinate respectively.

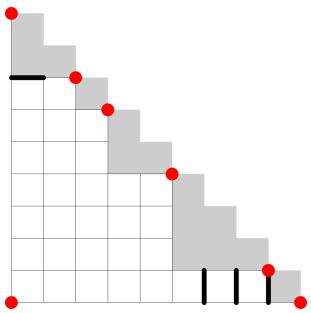

FIGURE 24. All length 1 chains.

We label the length 2 chains in Γ_λ as follows: let C_0 denote the length 2 chain near the origin. Let $C_2, C_3, \ldots, C_{2m-2}$ denote the length 2 chains that occur along the border of

Γ_λ where each even index chain C_{2k} (which may not exist) occurs near terminal vertex t_k and each odd index chain C_{2k-1} occurs at the corner of the kth shaded triangular subgrid corresponding to the coordinates that are fixed by k^{a_k} in λ. We call all odd index chains C_{2k-1}, $k = 2, \ldots, m-1$, the *type A* chains and all even index chains C_{2k}, $k = 1, \ldots, m-1$ (excluding C_0) the *type B* chains. As shown in Figure 25, the type B chain C_{2k} will not exist if $a_k = 1$ or $a_{k+1} = 1$ while all type A chains will exist. Our notation implicitly assumes that all of $C_2, \ldots, C_{2m-2}$ exist but all of our arguments hold with reference to the chains that actually exist for any given λ. For $0 \leq k \leq 2m-4$, let $X_{2,k}$ denote the set of facets in chains $C_0, C_2, C_3, \ldots, C_k$ and let $X_2 := X_{2,2m-2}$ denote the set of facets in chains $C_0, C_2, C_3, \ldots, C_{2m-2}$.

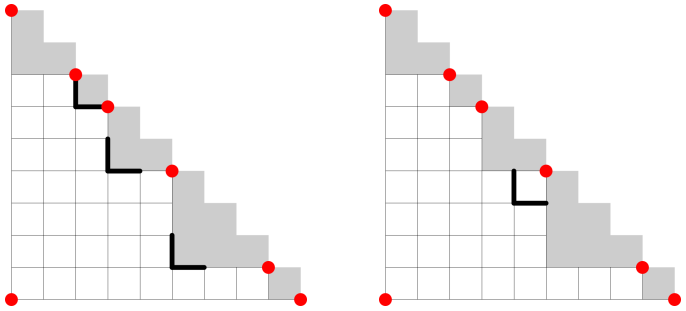

FIGURE 25. Left: all type A chains. Right: all type B chains.

Lemma 4.19. *The chains D_1, D_2, C_2, C_3, ... , C_{2m-2} can be ordered into a sequence D_1, C_2, C_3, ... , C_{2m-2}, D_2 such that any $\phi \in Aut(GT_\lambda)$ preserves the ordering within this sequence. In particular, any ϕ sends each chain in this sequence to itself or ϕ flips the order of the sequence.*

Figure 26 shows a picture of the sequence of chains mentioned in Lemma 4.19.

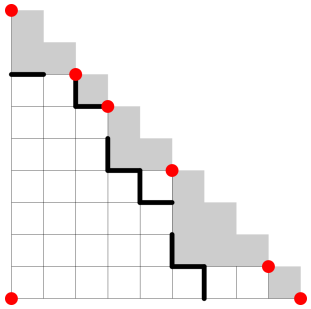

FIGURE 26. The sequence of chains $D_1, C_2, \ldots, C_{2m-2}, D_2$.

Proof. By Lemma 4.13, ϕ preserves lengths of chains. We say two length 1 chains are incompatible if there does not exist a vertex containing both of them and we say they are compatible otherwise. Note that all chains in $\mathcal{C}_1$ are incompatible with each other but compatible with any chain in $\mathcal{C}_2$. The set of chains that is the image of $\mathcal{C}_1$ under ϕ must all be incompatible with each other, so ϕ must map all chains of $\mathcal{C}_1$ to chains in $\mathcal{C}_1$ or it must all chains of $\mathcal{C}_1$ to chains in $\mathcal{C}_2$. The latter can only occur if the sizes of $\mathcal{C}_1$ and $\mathcal{C}_2$ are the

same. The rest of this proof and the proof of Lemma 4.20 will show that this can only occur if $\lambda = \lambda'$.

Thus $|O_G(X_1)| \leq (a_2)!^{\delta_{a_1,1}} \cdot (a_{m-1})!^{\delta_{a_m,1}} \cdot 2^t$ where $t = 1$ if $\lambda = \lambda'$ and $t = 0$ otherwise. Any $\phi \in G(X_1)$ must send C_0 to itself because it is the only length 2 chain such that for either of its facets, there exist vertices containing this facet and any facet in a type A chain. Note that ϕ may send the facets of C_0 to each other.

Chain D_i and a type A chain are incompatible if there does not exist a vertex containing the facet in the length 1 chain and both facets of the length 2 chain. Visually, the type A chain closest to the length 1 chain will be the only type A chain incompatible with it. Two type A chains are incompatible if there does not exist a vertex containing all four facets in these chains. Visually, two type A chains are incompatible iff they occur at the corners of adjacent shaded triangular subgrids (these are the subgrids corresponding to the fixed entries of $\mathcal{I}_n$). We can form a sequence starting with D_1 where two adjacent chains in this sequence are incompatible. Visually, this sequence corresponds to reading D_1, D_2 and the type A chains from left to right.

Now we insert the type B chains into this sequence. A type B chain and D_i are incompatible if there does not exist a vertex containing both facets in the type B chain and the facet in D_i. A type B chain and type A chain are incompatible if there does not exist a vertex containing all four facets in these chains. Visually, a type B chain is incompatible only with the chains adjacent to it. So we can insert the type B chains into the sequence so that the sequence corresponds to reading D_1, D_2 and all length 2 chains from left to right.

Note that $\phi \in G(X_1)$ must preserve these relations of incompatibility so ϕ fixes all length 2 chains. Finally, note that ϕ cannot map the two facets in a type A chain to each other: once D_1 is fixed, there does not exist a vertex containing the facet in D_1 and the vertical facet in the type A chain incompatible with D_1, but there does exist such a vertex for the horizontal facet. This fixes the image of the facets in this type A chain. Similarly, if the facets in a type A chain are fixed, then the facets in a type A chain incompatible with it will also be fixed. Of all length 2 chains, ϕ can only swap the facets in type B chains and chain C_0 and hence $|O_{G(X_1)}(X_2)| \leq 2^{r_1+1}$ where r_1 is the number of type B chains. □

Lemma 4.20. *If the sequence of chains $D_1, C_2, C_3, \ldots, C_{2m-2}, D_2$ is flipped by some $\phi \in Aut(GT_\lambda)$, then $\lambda = \lambda'$.*

Proof. Assume this sequence of chains is flipped. Define the distance between D_1 and C_3 (which always exists) to be the cardinality of the minimal set of facets $\mathcal{S}$, containing no other facets from this sequence, such that the intersection of facet D_1, one facet from C_3, and the facets in $\mathcal{S}$ is empty. We give examples of such sets in Figures 27. This distance is $a_1 + a_2 - 2$ since such a set must visually separate the coordinates that are fixed as 1 and 2 (the top-leftmost gray squares in Fig. 27). Similarly, the distance between D_2 and C_{2m-3} is $a_{m-1} + a_m - 2$. Note that ϕ must preserve this distance, as this property can be realized in the face lattice of the polytope. So if ϕ flips the sequence of chains, then $a_1 + a_2 = a_{m-1} + a_m$.

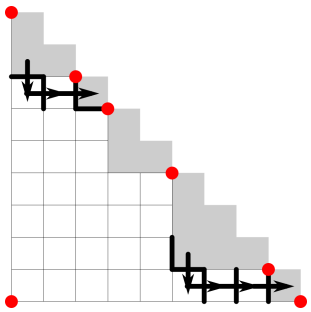

FIGURE 27. The distance between D_1 and C_3 is 1 and the distance between D_2 and C_{2m-3} is 2.

Recalling the definition of $\mathcal{G}_\lambda$, we can define the depth of a node as one plus the number of edges between it and the root. Then the depth of C_{2k-1} is

$$a_1 + \ldots + a_{k-1} + a_{k+1} + \ldots + a_m = \sum_{i=1}^{m} a_i - a_k.$$

Note that ϕ must send type A chains to type A chains. Since ϕ is an automorphism of $\mathcal{G}_\lambda$, it must preserve depth so if ϕ flips the sequence of chains, then it sends C_{2k-1} to chain $C_{2m-2k+1}$ so $a_k = a_{m+1-k}$ for $2 \le k \le m-1$. Since $a_1 + a_2 = a_{m-1} + a_m$, we have shown $a_k = a_{m+1-k}$ for $1 \le k \le m$ so $\lambda = \lambda'$. □

Now we are in a position to finish the proof of Theorem 1.3.

Proof of Thm. 1.3 cont. In Section 4.1, we showed that $(S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}) \ltimes_\varphi \mathbb{Z}_2^{r_2} \subseteq \mathrm{Aut}(\mathrm{GT}_\lambda)$. Now we show that the order of $\mathrm{Aut}(\mathrm{GT}_\lambda)$ is at most the order of this group.

By Lemmas 4.19 and 4.20, it suffices to bound $|G(X_1 \cup X_2)|$. Consider any automorphism $\phi \in G(X_1 \cup X_2)$. We can argue up the rooted tree $\mathcal{G}_\lambda$ to show that all of its vertices are fixed under ϕ. Specifically, ϕ must preserve adjacencies between chains so it is an isomorphism of the tree fixing the leaves of $\mathcal{G}_\lambda$ and hence must act trivially on the nodes. Furthermore, Lemma 4.18 implies that the orientation of any two non-leaves must be the same. If $\mathcal{G}_\lambda$ has at least two leaves, then the common ancestor of these two leaves cannot be flipped by ϕ or else the direct children of this node cannot be fixed. Otherwise, $\mathcal{G}_\lambda$ is a chain of nodes. If $m \ge 4$, then there are always at least two type A chains and hence at least 2 leaves. Similarly, if $m = 3$ and either $a_1, a_2 \ge 2$ or $a_2, a_3 \ge 2$, then there will be one type A chain and at least one type B chain. So $\mathcal{G}_\lambda$ can only be a chain of nodes if $\lambda = (1, 2^{a_2}, 3)$ or $\lambda = (1^{a_1}, 2, 3^{a_3})$. In the first case, $\mathcal{G}_\lambda$ contains two nodes, both of which have length 2. In the second case, WLOG assume that $a_1 \ge 2$. Consider the chain ending at a facet F that shares a vertex of Γ_λ with facet in D_1. An example of this chain is shown in Figure 28. If this chain were flipped, then there exists a vertex containing the image of F and the facet in D_1, but there does not exist a vertex containing F and the facet in D_1. Hence all of the non-leaves in $\mathcal{G}_\lambda$ are not flipped.

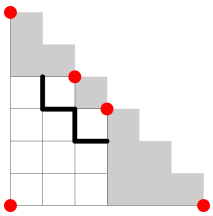

FIGURE 28. An example of the chain containing facet F.

Therefore, $|G(X_1 \cup X_2)| = 1$ and we have

$$\begin{aligned} |G| &= |O_G(X_1)||O_{G(X_1)}(X_2)||G(X_1 \cup X_2)| \\ &\leq a_2!^{\delta_{a_1,1}} \cdot a_{m-1}!^{\delta_{a_m,1}} \cdot 2^{r_1+1} \cdot 2^{r_2} \\ &= |(S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}) \ltimes_\varphi \mathbb{Z}_2^{r_2}|. \end{aligned}$$

Since $(S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}) \ltimes_\varphi \mathbb{Z}_2^{r_2}$ is contained in $\mathrm{Aut}(\mathrm{GT}_\lambda)$, we have $\mathrm{Aut}(\mathrm{GT}_\lambda) \cong (S_{a_2}^{\delta_{1,a_1}} \times S_{a_{m-1}}^{\delta_{1,a_m}} \times \mathbb{Z}_2^{r_1+1}) \ltimes_\varphi \mathbb{Z}_2^{r_2}$, as desired. □

ACKNOWLEDGEMENTS

This research was carried out as part of the 2016 REU program at the School of Mathematics at University of Minnesota, Twin Cities, and was supported by NSF RTG grant DMS-1148634 and NSF grant DMS-1351590. The authors are especially grateful to Victor Reiner for his mentorship and support, and for many fruitful conversations. The authors would also like to thank Elise delMas and Craig Corsi for their valuable advice and comments.

REFERENCES

[ABS11] Federico Ardila, Thomas Bliem, and Dido Salazar. Gelfand–Tsetlin polytopes and Feigin–Fourier–Littelmann–Vinberg polytopes as marked poset polytopes. *Journal of Combinatorial Theory, Series A*, 118(8):2454–2462, 2011.

[ACK16] Byunghee An, Yunhyung Cho, and Jang Soo Kim. On the f-vectors of Gelfand-Cetlin Polytopes. *arXiv preprint arXiv:1606.05957*, 2016.

[KST12] Valentina A Kirichenko, Evgeny Yu Smirnov, and Vladlen A Timorin. Schubert calculus and Gelfand-Zetlin polytopes. *Russian Mathematical Surveys*, 67(4):685, 2012.

[McA06] Tyrrell B McAllister. *Applications of Polyhedral Geometry to Computational Representation Theory*. PhD thesis, University of California Davis, 2006.

[Sta86] Richard P Stanley. Two poset polytopes. *Discrete & Computational Geometry*, 1(1):9–23, 1986.

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139

E-mail address: gaoyibo@mit.edu

Department of Mathematics, Yale University, New Haven, CT 06520

E-mail address: benjamin.krakoff@yale.edu

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139

E-mail address: lisayang@mit.edu